

METHODS AND SYSTEMS FOR REGULATING RESOURCE USAGE

RELATED APPLICATIONS

[0001] This application relates to Attorney Docket No. GP-175-31-US, filed herewith and entitled, "Systems and Methods for Providing and Installing Software," the entirety of which is incorporated herein by reference. This application also relates to Attorney Docket No. GP-175-32-US, filed herewith and entitled, "Systems and Methods for Providing Software Updates," the entirety of which is incorporated herein by reference.

FIELD OF THE INVENTION

[0002] The invention generally relates to methods of regulating resource usage of an application.

BACKGROUND

[0003] As World Wide Web ("web") search engines (sometimes referred to as "Internet Search Engines") have improved, many users have turned to these search engines for navigating the web, rather than inputting uniform resource locators (URLs) into browser address fields or using browser bookmarks. Search engines may perform searches of various databases, which may be public, e.g., the Internet, and/or private, e.g., an intranet, a client device, etc., using one or more known search techniques. For example, one known search technique, described in an article entitled, "The Anatomy of a Large-Scale Hypertextual Search Engine," by Sergey Brin and Lawrence Page, assigns a degree of importance to a document, such as a web page, based on the link structure of the web.

[0004] These search engines may be applied to information over the web or over information on a user's local hard drive. Some search methods often involve computationally intensive applications. Some of these applications may fail as a practical matter, e.g., resulting in the user interrupting or discontinuing use of the system, because of the demands they place on client-device resources. Such resource demands can slow response time of the client device, use excessive amounts of central processing unit space, and use excessive amounts of processor memory, such as RAM, which often results in file swapping, and can degrade the overall user experience in certain circumstances.

[0005] What is needed is an improved method and system – especially in a high-bandwidth server environment – to dynamically regulate, e.g., throttle, the application's usage of the client device such that the user's perception of system performance is not affected in a material way by interaction with the application.

SUMMARY OF THE INVENTION

[0006] Embodiments of the present invention comprise methods and systems for regulating resource usage. In one exemplary embodiment, a method comprises receiving an operating parameter of the client device, assigning a value to a usage variable associated with the operating parameter of the client device, and correlating by an application a resource usage level of the application with the usage variable.

[0007] Another exemplary embodiment comprises a computer readable medium comprising instructions, that, when executed, cause an application to perform the steps of receiving an operating parameter of a client device, assigning a value to a

usage variable associated with the operating parameter of the client device, and correlating a resource usage level of the application with the usage variable.

[0008] In both embodiments above, multiple operating parameters can be used simultaneously to generate a “resource utilization measure” which can be assigned to the usage variable and correlated with the resource usage level.

[0009] These exemplary embodiments are mentioned not to limit or define the invention, but to provide examples of embodiments of the invention to aid understanding thereof. Exemplary embodiments are discussed in the Detailed Description, and further description of the invention is provided there. Advantages offered by the various embodiments of the present invention may be understood by examining this specification.

BRIEF DESCRIPTION OF THE DRAWINGS

[0010] The accompanying drawings, which constitute part of this specification, help to illustrate embodiments of the invention.

[0011] Figure 1 is a block diagram of an exemplary environment for implementing an embodiment of the present invention.

[0012] Figure 2 is a block diagram of a method according to an embodiment of the present invention.

DETAILED DESCRIPTION

[0013] Embodiments of the present invention include methods and systems for regulating resource usage. Referring now to the drawings in which like numerals indicate like elements throughout the several figures, Figure 1 is a block diagram illustrating an exemplary environment for implementation of an embodiment of the

present invention. The system 100 shown in Figure 1 includes multiple client devices 102a-n in communication with a server device 104 over a network 106. The network 106 shown comprises the Internet. In other embodiments, other networks, such as an intranet may be used. Moreover, methods according to the present invention may operate within a single computer.

[0014] The client devices 102a-n shown each include a computer-readable medium, such as a random access memory (RAM) 108, coupled to a processor 110. As will be described in further detail below, the client devices 102a-n each include a resource regulating application. The processor 110 executes computer-executable program instructions stored in memory 108. Such processors may include a microprocessor, an ASIC, and state machines. Such processors include, or may be in communication with, media, for example computer-readable media, which stores instructions that, when executed by the processor, cause the processor to perform the steps described herein. Embodiments of computer-readable media include, but are not limited to, an electronic, optical, magnetic, or other storage or transmission device capable of providing a processor, such as the processor 110 of client 102a, with computer-readable instructions. Other examples of suitable media include, but are not limited to, a floppy disk, CD-ROM, DVD, magnetic disk, memory chip, ROM, RAM, an ASIC, a configured processor, all optical media, all magnetic tape or other magnetic media, or any other medium from which a computer processor can read instructions. Also, various other forms of computer-readable media may transmit or carry instructions to a computer, including a router, private or public network, or other transmission device or channel, both wired and wireless. The instructions may

comprise code from any computer-programming language, including, for example, C, C++, C#, Visual Basic, Java, Python, Perl, and JavaScript.

[0015] Client devices 102a-n may also include a number of external or internal devices such as a mouse, a CD-ROM, DVD, a keyboard, a display, or other input or output devices. Examples of client devices 102a-n are personal computers, digital assistants, personal digital assistants, cellular phones, mobile phones, smart phones, pagers, digital tablets, laptop computers, Internet appliances, and other processor-based devices. In general, a client device 102a may be any type of processor-based platform that is connected to a network 106 and that interacts with one or more application programs. Client devices 102a-n may operate on any operating system capable of supporting a browser or browser-enabled application, such as Microsoft® Windows® or Linux. The client devices 102a-n shown include, for example, personal computers executing a browser application program such as Microsoft Corporation's Internet Explorer™, Netscape Communication Corporation's Netscape Navigator™, and Apple Computer, Inc.'s Safari™.

[0016] Through the client devices 102a-n, users 112a-n can communicate over the network 106 with each other and with other systems and devices coupled to the network 106. As shown in Figure 1, a server device 104 is also coupled to the network 106. The server device comprises a server processor 116 and a server memory storage device 118. The server device 104 shown comprises a single computer. However, in other embodiments, the server device 104 may comprise multiple and/or distributed devices, or there may be no server device.

[0017] In one embodiment, a user 112a-n generates a search query 114 at a client device 102a. The client device 102a transmits the query 114 to the server device 104 via the network 106. For example, a user 112a types a textual search query into a query field of a web page of a search engine interface or other client-side software displayed on the client device 102a, which is then transmitted via the network 106 to the server device 104.

[0018] In the embodiment shown, a user 112a inputs a search query 114 at a client device 102a, which transmits an associated search query signal 122 reflecting the search query 114 to the server device 104. The search query 114 may be transmitted directly to the server device 104 as shown. In another embodiment, the query signal 122 may instead be sent to a proxy server (not shown), which then transmits the query signal 122 to server device 104. Other configurations are possible.

[0019] The server device 104 shown includes a server executing a search engine application program, such as the Google™ search engine. Similar to the client devices 102a-n, the server device 104 shown includes a processor 116 coupled to a computer-readable memory 118. Server device 104, depicted as a single computer system, may be implemented as a network of computer processors. Examples of a server device 104 are servers, mainframe computers, networked computers, a processor-based device, and similar types of systems and devices. Client processor 110 and the server processor 116 can be any of a number of computer processors, such as processors from Intel Corporation of Santa Clara, California and Motorola Corporation of Schaumburg, Illinois.

[0020] Memory 118 contains the search engine application program, also known as a search engine 120. The search engine 120 locates relevant information in response to a search query 114 from a user 112a-n. The search engine 120 then provides the result set 124 to the client 102a via the network 106.

[0021] In the embodiment shown, the server device 104, or related device, has previously performed a crawl of the network 106 to locate articles, such as web pages, stored at other devices or systems connected to the network 106, and indexed the articles in memory 118 or on another data storage device. Articles include, for example, web pages of various formats, such as HTML, XML, XHTML, Portable Document Format (PDF) files, and word processor, database, and application program document files, audio, video, or any other documents or information of any type whatsoever made available on a network (such as the Internet), a personal computer, or other computing or storage means. The embodiments described herein are described generally in relation to HTML files or documents, but embodiments may operate on any type of article, including any type of image.

[0022] It should be noted that various embodiments of the present invention may comprise systems having different architecture than that which is shown in Figure 1. For example, in some systems according to the present invention, server device 104 may comprise a single physical or logical server, or there may be no server. The system 100 shown in Figure 1 is merely exemplary, and is used to explain the exemplary method shown in Figure 2.

[0023] Various methods may be implemented in the environment shown in Figure 1 and other environments according to the present invention. For example, in one

embodiment, a resource regulating application 109 receives information regarding utilization of a processor 110 of the client device 102a-n. At a predetermined percentage of processor 110 utilization, the application 109 scales back (i.e., throttles) its requirements for processor 110 usage.

[0024] Thus, if the processor 110 is heavily utilized, the application 109 throttles-back its processor 110 requirements. This method alleviates performance degradation of the client device 102a-n that would be observable by the user 112a-n of the client device 102a-n if the application 109 utilized resources without regard to existing demands on the processor 110.

[0025] Throttling-back processor 110 requirements can be accomplished by pausing between execution or processor-intensive calculations, such as might be performed for some indexing steps. A similar process may occur for other resources, such as for example, memory 108 resources and communication resources (not shown).

Throttling-back memory 108 can be accomplished by loading fewer components at one time or by having each component dynamically scale back on its memory usage. Communication (or bandwidth) requirements can be throttled-back by slowing download from the server 104, e.g., setting a lower priority or requesting less data at any one time. Of course, the reverse can be performed, i.e., throttling-up when resources become available.

[0026] Execution of the application may only comprise execution of a step or sequence of steps before again checking the processor 110 utilization. The application 109 may include logic that only allows a small sequence of steps to execute before checking a parameter that specifies whether throttling should be on or

off. Thus, the steps above may be repeated almost constantly while the application 109 is active, which permits dynamic throttling in response to changing resource availability.

[0027] Of course, third party components (not shown) may interact with the system 100. For example, resource demand is called upon only when performance measures indicate that reasonable resources exist to execute the application 109.

[0028] Referring now to Figure 2, a method 200 according to an embodiment of the present invention is shown. The method 200 may be employed to regulate usage of a client device, such as for example, in the environment of the system 100 described above. Items shown in Figure 1 are referred to in describing method 200 to aid understanding of this embodiment. However, methods according to other embodiments of the present invention may be employed in a wide variety of other environments.

[0029] As indicated by block 210, the method 200 comprises receiving one or more operating parameters of the client device 102a. Examples of operating parameters may include processor usage or load, number of processes and threads running, processor speed, memory capacity or size, power consumption, virtual memory paging, disk size, free space, transfer rate, I/O operations per second, and communication bandwidth, among others.

[0030] In one embodiment, receiving the operating parameter comprises monitoring the operating parameter. The operating parameter can be monitored by the application 109, by the client device 102a itself, or by the server 104. In this embodiment, the monitored parameter can then be transmitted to the application 109.

[0031] In one embodiment, the operating parameter of the client device comprises a client processor load. Generally, the load of the client device processor 110 comprises an amount of a total processing capacity of the client device processor 110 being used. Thus, the load is typically expressed in terms of a percentage of the total capacity. For example, a processor load of 50% means that, over some period of time, one-half of the processor 110 is being used, while the other half of the processor 110 is still available for use.

[0032] In one embodiment, receiving the operating parameter of the client device 102a, e.g., calling specific performance information, can be achieved by the resource regulating application 109 transmitting an inquiry signal (not shown) to the client device 102a. In one embodiment, the inquiry signal comprises code or data structure that causes the processor 110 to obtain the desired information and transmit that information to the application 109. Alternatively, receiving the operating parameter can be accomplished by other suitable means, such as active reporting by the processor 110 to the application 109, or calling an operating system function or set of functions.

[0033] In another embodiment, receiving the operating parameter of the client device 102a can be achieved by the server device 104 communicating an inquiry signal (not shown) to the client device 102a. In one embodiment, the inquiry signal comprises code or data structure that causes the client device 102a to obtain the desired information and transmit that information to the server device 104. Alternatively, receiving the operating parameter can be accomplished by other

suitable means, such as active reporting by the client device 102a to the server device 104.

[0034] In another implementation of receiving 210, receiving the parameter comprises receiving the parameter during an initial load of the client processor 110. For example, the initial load may include the amount of a total capacity of the client processor 110 being used when the client device 102a begins operation or when communication is initiated or established between the client device 102a and server device 104 or between the client device 102a and the application 109. An initial load of 50% indicates that half of the computational capacity of the client processor 110 is available when the load is first received – usually when the client device 102a begins operation or when it first begins communicating with the server device 104 or with the application 109.

[0035] In one embodiment, receiving the parameter comprises receiving the parameter during a predetermined time interval. Suitable predetermined time intervals can be used, such as for example, five seconds. Alternatively, the parameter can be received substantially continuously. In an alternate embodiment, the operating parameter is received during an initial load and at the predetermined time intervals while the server device 104 or the application 109 and the client device 102a are in communication with one another.

[0036] In another implementation of receiving 210, receiving 210 further comprises monitoring a period of inactivity of the client device 102a. In one embodiment, monitoring comprises monitoring the inactivity of the client processor 110. Alternatively, the inactivity of other features of the client device 102a, such as for

example, memory usage, event queue usage, and event frequency can be monitored.

In an embodiment, the period of inactivity comprises a first time and a second time.

In one embodiment, the second time is greater than the first time.

[0037] In another implementation of receiving 210, receiving the operating parameter further comprises receiving information regarding a bandwidth use of the client device 102a. The bandwidth may include, for example, the total capacity of the available bandwidth, an amount of bandwidth being used at a given time, and a frequency of bandwidth usage. These receiving embodiments discussed above are merely examples. One or more of these or other operating parameters may be received and/or monitored and used.

[0038] Bandwidth generally refers to an amount of information that can be carried or transported over a wired or wireless communication pathway or network in a given period of time. A broadband bandwidth is generally one that may be able to carry sufficient information to transmit and sustain a relatively bandwidth-intensive application, such as a succession of images in a video conference or presentation. Bandwidth otherwise may refer to the width of the range of frequencies that an electronic signal occupies on a transmission medium. Thus, any digital or analog signal has a bandwidth.

[0039] In a digital system, including the system 100 described in Figure 1, bandwidth is expressed as bits of data per second ("bps"). Thus, a device, such as a modem, that works at 57,600 bps has twice the bandwidth of a device that works at 28,800 bps. In analog systems, bandwidth is expressed in terms of the difference between the highest-frequency signal component and the lowest-frequency signal

component (generally measured in hertz). Communications pathways usually comprise a succession of links, each with its own bandwidth. If one of these links is slower than the rest, it is generally referred to as a bandwidth bottleneck.

[0040] The various devices of the system 100 shown in Figure 1, may each comprise a link in a communication pathway. As described above, the multiple client devices 102a-n are in communication with the server device 104 over network 106. The network 106 comprises the Internet, but in other embodiments, the network may comprise an intranet. The devices 102a-n may communicate with the network 106 via modems (not shown) of various bandwidths, such as 57.6 kbps or broadband. The modems of the devices 102a-n may communicate with the network 106 over telephone lines, fiber-optic cables, co-axial cable, or wirelessly, such as with routers, repeaters, or via satellite. Other suitable means of communication can be used.

[0041] As will be discussed below, resource throttling can be modified according to varying, predetermined periods of inactivity. For example, a given resource usage will be allowed at the first time, e.g., three minutes. However, at a second time, e.g., five minutes of inactivity, the resource usage can be throttled-up. Likewise, at a third time, e.g., two minutes, the resource usage of the client device 102a can be throttled-back. In one embodiment, a time less than the first time will prohibit a resource usage, e.g., throttling resource usage to zero.

[0042] As indicated by block 220, the method 200 further comprises assigning a value to a usage variable associated with the operating parameter of the client device 102a. In one embodiment, the assigning step 220 is performed by the application

109. In other embodiments, the assigning step 220 is performed by the server device 104, the client processor 110, or other device or processor.

[0043] In one embodiment, the usage variable is a quantitative performance measure of the client device 102a associated with (or based at least in part on) the operating parameter. Usage variables express a measure or performance characteristic of a specific operating parameter, and may be equal to the operating parameter. For example, the operating parameter received is a load of the client processor 110. A usage variable of 25% load indicates that three-quarters of the available computational capacity of the client processor 110 is available for usage.

[0044] In one embodiment, the value is an ordinal ranking of the usage variable representing a qualitative performance measure of the client device 102a. For example, the qualitative performance measure can be expressed as low, medium, and high. Other values may be assigned, as these are only examples.

[0045] As one example, when the client device 102a has a load of 25%, it likely will perform better (i.e., the user 112a will not likely perceive that its interaction with the client device 102a has slow response time or is otherwise affected) than when it has a load of 75%. Thus, one can assign a value of "1" where the load of the processor 110 is 25% or less, a value of "2" where the load is between 25 and 50%, a value of "3" where the load is between 50% and 75%, and a value of "4" where the load is greater than 75%. In such a scale, the value of "1" represents the greatest predicted level of performance for the client device 102a and the value of "4" represents the lowest predicted level of performance of the client device 102a.

[0046] Other suitable values and usage variables can be used. For example, the scale of values can include more than the four gradations described above. The values assigned to usage variables will depend on various factors, including the type of operating parameter(s) selected for monitoring.

[0047] As another example, one can assign a usage variable to a bandwidth parameter, such as for example, bandwidth load. The total bandwidth may include for example, the bandwidth available over devices such, as a modem for communication between the client device 102a and the server. Thus, a bandwidth load of 25% would indicate that one-quarter of the available bandwidth is being used, leaving the remaining three-quarters of the total bandwidth as available. Thus, one can assign a value of "1" where the bandwidth load is 25% or less, a value of "2" where the load is between 25% and 50%, a value of "3" where the load is between 50% and 75%, and a value of "4" where the load is greater than 75%. In such a scale, the value of "1" represents the greatest predicted level of performance for the client device 102a and the value of "4" represents the lowest predicted level of performance of the client device 102a. Other suitable usage variables and values can be used.

[0048] As indicated by block 230, the method 200 further comprises correlating by the application 109 a resource usage level of the application 109 with the usage variable. In correlating the resource usage level, the application 109 modifies its own execution based on the usage variable. In this manner, the resource usage can be regulated or "throttled." Several examples follow.

[0049] The following table illustrates an example of one such embodiment. In the table, the operating parameter is the usage of the client processor 110 and the usage

variable is expressed as the load of the client processor 110. For each usage variable, there is a corresponding resource usage level (right column) by the application 109.

Client Processor Load	Resource Usage Level
< 25%	15%
25% to 50%	10%
50% to 75%	5%
> 75%	1%

[0050] As seen in the table above, the resource usage level corresponds to the level of “throttling.” Thus, where the client processor load is less than or equal to 25%, the resource usage is fully throttled. Where the client processor load is greater than 75%, the resource usage is throttled back. In this example, a client processor load of 25% correlates to a resource usage level by the application 109 of approximately 15%, whereas a client processor load of greater than 75% correlates to a resource usage level by the application of approximately 1%. The resource usage may also change smoothly with changes in the usage variable. For example, if the usage variable indicates the percentage of processor load, the resource usage may be set to $(100 - \text{Current Load})/2$, i.e., one-half of the available load.

[0051] The amount of usage, i.e., the resource usage level, by the application 109 in this exemplary embodiment is further correlated with the period of inactivity of the client device 102a, i.e., the first time. In this exemplary embodiment, the period of inactivity is less than approximately three minutes. As shown in the table above, where the period of inactivity of the client device 102a is less than approximately

three minutes and the client processor load is less than 25%, the resource usage level of the client processor 110 is limited to approximately 10%.

[0052] In another embodiment, the method 200 further comprises correlating the resource usage level of the application 109 with the second time, i.e., greater than approximately three minutes. Although not shown in the table above, the resource usage level of the application 109 at the second time is greater than that shown for the first time, i.e., three minutes, so long as the load of the client processor 110 is less than approximately 75%. Where the client processor load is greater than approximately 75%, the resource usage level of the application 109 is generally limited to approximately 5%.

[0053] For example, where the period of inactivity of the client device 102a is greater than three minutes and where the processor load is less than 25%, the resource usage level of the application 109 is greater than 15%. Other suitable parameters and correlations can be used.

[0054] As discussed above, the parameters are received and/or monitored in the embodiment shown at a predetermined interval, and the correlation is adjusted with changing values over time. Likewise, the period of inactivity of the client device 102a is received and/or monitored at a predetermined interval. In one embodiment, the inactivity of the client device 102a is monitored substantially continuously.

[0055] The following table illustrates another embodiment of correlation 230. In this alternate embodiment, the parameter comprises a first parameter and a second parameter. The first parameter comprises a speed of the client processor 110 and the second parameter comprises a capacity of the client memory storage device 108.

Thus, the resource usage level of the application 109 can be a combination of two or more parameters.

Client Device Performance	Client Processor Speed	Capacity of Client Memory Storage	Resource Usage Level
1	300 MHz to 800 GHz	64 MB to 128 MB	5 MB
2	800 MHz to 1.5 GHz	128 MB to 255 MB	5 MB to 10 MB
3	> 1.6 GHz	129 MB to 256 MB	10 MB to 20 MB

[0056] The first column indicates the value, expressed in an ordinal ranking of a measure of predicted client device 102a performance, where 1 represents “low” performance, 2 represents “medium” performance, and 3 represents “high” performance. As discussed above, performance levels represent expected or predicted levels of client device 102a performance.

[0057] Thus, a client device 102a having a “high” performance is expected to be able to devote more of its resources to the application 109 than a client device 102a having a “low” performance and still be able to carry out its other tasks without the user 112a observing readily noticeable or degraded performance, such as response time, of the client device 102a.

[0058] The second column includes the first parameter, which comprises the speed of the client processor 110, expressed in megahertz (MHz) and gigahertz (GHz). The third column includes the second parameter, which comprises the capacity of the client memory storage device 108, expressed in megabytes (MB). The fourth column is the maximum resource usage of the application 109, also expressed in megabytes.

[0059] The following table illustrates another embodiment of correlation 230. In this alternate embodiment, the operating parameter comprises a bandwidth load. In the table, the parameter is the usage of the bandwidth of the client device 102a and the usage variable is expressed as the bandwidth load. For each value, there is a corresponding usage level (right column) of the bandwidth of the client device 102a.

Client Bandwidth Load	Resource Usage Level
< 25%	15%
25% to 50%	10%
50% to 75%	5%
> 75%	1%

[0060] As seen in the table above, the usage level of the client bandwidth corresponds to the level of “throttling.” Thus, where the client bandwidth load is less than or equal to 25%, the resource usage is fully throttled. Where the client bandwidth load is greater than 75%, the resource usage is throttled back. In this example, a client bandwidth load of 25% correlates to a resource usage level of approximately 15%, whereas a client bandwidth load of greater than 75% correlates to a resource usage level of approximately 1%.

[0061] As discussed above, the parameters are received and/or monitored at an initial load, and at a predetermined interval. The correlation is adjusted with changing usage variables. The correlations described above generally reduce the likelihood of occurrences of file swapping by the client device 102a. File swapping typically occurs when the client device 102a lacks adequate memory, which results in

delay while the client device 102a loads the swapped application when sufficient memory 108 becomes available. Other suitable parameters and correlations can be used.

[0062] In another embodiment, the application 109 uses the client device 102a as the client device 102a writes to the computer readable medium of the client memory storage device 108. Preferably, this is done such that the user 112a of the client device 102a will not likely perceive that its interaction with the server device 104 affects the client device's 102a response time or other capability.

[0063] In another embodiment, correlating by an application 109 a resource usage level of the application 109 with the usage variable comprises performing an activity affecting a usage variable proximate to a time that the usage variable indicates an existing activity. For example, the application 109 detects when other applications or the operating system of the client device 102a is writing to the computer local storage device, e.g., hard drive, ZIP drive, etc. It then performs its own writing activity to the local storage device at the same time. This can result in multiple improvements, for example combining the activity with existing activity can reduce the overall work required when writing to disk, and the user may prefer to have disk activity combined with existing disk activity than introduce new periods of disk activity which may be distracting. The method is particularly beneficial for mobile computing devices, such as laptops, where power conservation is very important, and the disk may stop spinning to reduce power usage. Further, reading or writing to disk can cause the operating system to swap inactive memory from other applications to disk, therefore, if an application reads or writes to disk during periods of inactivity, the other

applications the user is using may be partially or completely swapped to disk when they return.

[0064] In another implementation of correlating 230, correlating by an application 109 a resource usage level of the application 109 with the usage variable comprises suspending one or more operations when the usage variable exceeds a threshold. In still another implementation of correlating 230, correlating by an application 109 a resource usage level of the application 109 with the usage variable comprises adjusting a rate of operation based at least in part on the usage variable.

[0065] In yet another implementation of correlating 230, correlating by an application 109 a resource usage level of the application 109 with the usage variable comprises adjusting a sequence of operations based at least in part on the usage variable. For example, operations may be performed in sequence rather than in parallel. In still a further implementation of correlating 230, correlating by an application 109 a resource usage level of the application 109 with the usage variable comprises adjusting an active feature based at least in part on the usage variable. For example, a given feature may be turned off or disabled if there is insufficient capacity in memory 108.

[0066] A computer readable medium of a server device, a resource regulating application, processor, or other device or application comprises instructions, that, when executed, cause the server device, application, processor, or other device or application to perform method 200. Preferably, the server device, resource regulating application, and the computer readable medium are similar to that described above and with reference to Figure 1. Alternatively, other suitable server devices,

applications, computer readable media, processors, or other devices or applications can be used. The configuration of the system 100 and the method 200 will not be repeated here.

[0067] While the present invention has been disclosed with reference to certain embodiments, numerous modifications, alterations, and changes to the described embodiments are possible without departing from the sphere and scope of the present invention, as defined by the appended claims. Accordingly, it is intended that the present invention not be limited to the described embodiments, but that it has the full scope defined by the language of the following claims, and equivalents thereof.